

REMARKS

Explanation of Amendments

Independent claims 3 and 22 have been amended for clarification. No other claim has been amended, added, or deleted. Claims 3, 5-12, and 19-28 remain pending. Claims 3 and 22 are the independent claims.

Interview Summary

Application appreciates the time and attention of Examiner Simitoski during an interview with Mr. Robertson, Mr. Kamis, and Mr. Dunnam on September 15, 2009. During that interview, the invention was described in the context of the cited references. Examiner Simitoski indicated that the arguments needed better support in the claims and suggested that the independent claims be amended to more clearly identify the determination of the connection source that initiated the connection session. The amendments proposed above were submitted to Examiner Simitoski, who indicated that the proposed amendments would likely overcome the rejections but that the amended claims would likely require further consideration and search. Applicant has thus filed this response with an RCE. Further consideration is requested.

Claim Rejections – 35 USC §103

Claims 3, 6, 7, 9, 21-24, and 28 stand rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton (US 7,120,931) in view of Vaidya (US 6,279,113) and Jeffries (US 2003/0196095). Claims 5 and 8 stand rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton, Vaidya, and Jeffries in view of Desai et al. (US 2003/0188189). Claims 10, 12, and 25 stand rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton, Vaidya, and Jeffries in view of Daizo (US 6,424,654). Claim 11 stands rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton, Vaidya, Jeffries and Daizo in view of Desai et al. Claims 19 and 26 stand rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton, Vaidya, and Jeffries in view of Trcka et al. (US 6,453,345). Finally, claims 20 and 27 stand rejected under 35 USC §103(a) as allegedly being unpatentable as obvious over Cheriton, Vaidya, and Jeffries in view of Ullmann et al. (US 2002/0174362). In short, the examiner has repeated the previous rejections except that he has added Jeffries for a

purported teaching of detecting the spread of malicious software by counting references of destination addresses and comparing the count to a threshold to determine if a gateway may be receiving random addresses associated with malicious software. This is not what is claimed. The rejections are again traversed.

Repeated Arguments Regarding Cheriton and Vaidya

As previously noted, the claimed system and method detects surveillance probes on a computer communications network to alert sites to possible imminent intrusions by detecting the system surveillance (or scanning) that often precedes such attacks. For example, surveillance probes are used to gather information including whether hosts are alive, what services they are running, whether they are running a particular vulnerable application, etc. The claimed system and method exploits the fact that someone attempting to gather information about an enterprise to attack does not know what information they are attempting to gather, and because of that, they make mistakes. Almost all legitimate protocols on the IP Internet are query response type protocols, which means that both parties send packets to each other. Accordingly, the potential intruder must probe hosts in an attempt to gather information to determine when and where to attack the system.

The claimed system and method may exploit the fact that surveillance probes often receive no response or the response is an error message since the surveillance probes do not necessarily know what information they seek. In other words, by tracking the initiator of "bad" connections, the attacking source(s) may be identified. For example, when one IP source is responsible for too many unique bad connections within a specified time interval, an alert is raised. A "bad" connection may be defined very simply as any IP communication which does not generate any kind of valid reply. If one IP source racks up enough of these "bad" (or unreplied) connections to unique destinations, it is considered to be possibly surveying the network and an alert is raised. If it continues to send probes, further alerts will occur.

The claimed system and method accomplish this result by assembling data packets into "connection sessions" and deciding whether they are from "good" or "bad" sources. Assembling packets into connections is done by storing information about the connection in a memory so as to uniquely identify a particular connection. For example, the stored information may include statistics on source IP address that initiate probes over long periods

of time (*see e.g.* specification at paragraph [0051]). Applicant notes that the definitions of "source" and "destination" as used herein are from the perspective of the originator of the connection, where the source refers to information about the client and destination refers to information about the server. Thus, when a packet travels from the client to the server, the sources and destinations observed in the packet are the same as the connection sources and destinations. However, for packets traveling from the server to the client, the observed packet sources and destination are actually reversed.

The claimed system and method determines directionality by considering one packet in relation to any other packets which it has seen before for the same connection identification. In practice, the surveillance detector uses the first packet observed in a connection as an indication of the directionality of the connection. For example, as noted in paragraph [0070] of the specification, flag fields in the connection extrapolator may record from which direction interesting events, such as a first packet, initiate. In other words, the system and method uses timing-based directionality, meaning that the first packet determines the connection source. All subsequent packets match the existing entry in the connection table and may simply update any necessary information from that table.

In order to perform the necessary good/bad connection tracking (where, *e.g.*, "good" connections have bidirectional traffic and "bad" connections have unidirectional traffic), in addition to recording whether the low address is the connection source or destination, the system also tracks whether it has seen a connection from both the connection source and destination (or as illustrated in paragraphs [0069]-[0070] of the specification, whether the system has seen the connection from the low or high addresses). Other information, such as the number of packets observed, or the number of bytes of payload, and similar aggregate statistics about features of the packet can also be accumulated in the connection sessions. These aggregate statistics may be of use to the surveillance detector or users of the surveillance detection alerts. When sufficient bad evidence has built up regarding a connection, the system can generate an alert summarizing the bad activity.

These characteristic features of the method are set forth in claim 3. In particular, claims 3 recites a method of detecting surveillance probes on a computer communications network, comprising:

receiving a plurality of messages from a data sensor located at a network audit point, said data sensor sampling data packets on said computer communications network and outputting said messages, each of said messages describing an event occurring on said communications network;

processing said messages to form extrapolated connection sessions from said sampled data packets from which to determine a connection source that initiated the connection session by clustering packets exchanged between an address of said connection source and a destination address that are a) within a specified time period where the source and destination addresses are not predetermined, (b) have certain flags set, or c) have source and destination addresses that are not predetermined but have similar characteristics; and

detecting a surveillance probe by:
grouping said connection sessions into a plurality of groups of related connection source addresses;
scoring each group based on at least a quantity of attack destinations;
and
generating an alert for each group whose score is greater than an empirically derived threshold.

Similar features may be found in system claim 22. Such features are not shown or suggested by the cited prior art. In particular, Applicant submits that none of the cited art teaches “processing said messages to form extrapolated connection sessions from said sampled data packets from which to determine a connection source that initiated the connection session.” On the contrary, systems like that of Cheriton, which detects “flow data,” cannot perform this connection source identification since the received flow data does not have sufficiently precise information to establish who initiated the connection. Also, Cheriton does not detect a surveillance probe based on scoring whether or not the packets in response to a query indicate successful communication from the network protocol perspective or indicate attack destinations and then generate an alert if a threshold is exceeded as claimed.

As previously noted, Cheriton discloses a system and method for analyzing high speed data entering a router or firewall to identify detailed characteristics of the packets involved in an attack or a failure. The method includes generating filters based on analyzed flow data by separating the data into different network flows, analyzing at least one of the network flows, and detecting potentially harmful network flows. A filter is generated to prevent packets corresponding to the detected potentially harmful network flows from

passing through the network device. As illustrated in Figures 3 and 4, a netflow mechanism processes each packet in the network flow responsive to the entry for the network flow in the flow cache, and the netflow mechanism implements administrative policies that are designated for each network flow rather than for each packet. The network flows are analyzed and information on incoming packets is provided without examining each packet received. This “flow collection aggregation” allows for data to be stored by aggregate summary records instead of raw data records (column 6, lines 56-65). Once a group of packets is identified as harmful, the corresponding network flows may be analyzed to further refine the filter. The flow analyzer monitors the statistics associated with the aggregate filters and, if the statistics associated with an aggregate filter entry indicate a potential problem, creation of netflow entries is enabled for packets matching the entry. The flow analyzer receives a flow record for each flow matching the aggregate, and the flow generator determines how to refine the aggregate filter.

In contrast with the claimed method, Cheriton’s analysis of the flow data does not enable Cheriton’s system to “form extrapolated connection sessions from said sampled data packets from which to determine a connection source that initiated the connection session” as claimed. On the contrary, those skilled in the art will appreciate that the flow data of Cheriton are not connections and thus cannot form “connection sessions” as now claimed. The flow data analyzed by Cheriton is more like half connections since the flow data is extracted from a unidirectional traffic flow. Those skilled in the art will appreciate that the claimed “connection sessions” instead involve bidirectional flows of data where the connection source (client) and connection destination (server) are known. Since the flow data from a router/switch analyzed by Cheriton does not include sufficient information to permit one to accurately distinguish the connection client and server by, for example, assembling flow data into “connection session,” and since Cheriton provides no teachings regarding “determining a connection source that initiated the connection session,” the claimed method and system are not suggested by Cheriton. Moreover, since Cheriton does not sample the network data but instead processes the entire network flow, no “extrapolated connection sessions” are generated as claimed.

Furthermore, as acknowledged by the Examiner, Cheriton does not score each group or generate an alert for each group whose score is greater than an empirically derived

threshold. For a teaching of scoring each group and generating an alert for each group whose score is greater than an empirically derived threshold, the Examiner references Vaidya's teachings of counting characteristics in the packet stream, such as an attempt to access a file, and determining whether the count exceeds a threshold. However, Vaidya does not teach detecting a surveillance probe by grouping connection sessions into a plurality of groups of related connection source addresses and scoring each group "based on at least a quantity of attack destinations" as now claimed. Applicant further submits that Vaidya does not teach the above-mentioned features that are not taught by Cheriton. Accordingly, even if the teachings of Vaidya could have been combined with the teachings of Cheriton as the Examiner alleges, the claimed system and method would not have resulted.

Jeffries Does Not Overcome Shortcomings of Cheriton and Vaidya

The examiner has repeated the previous rejections except that he has added Jeffries for a purported teaching of detecting the spread of malicious software by counting references of destination addresses and comparing the count to a threshold to determine if a gateway may be receiving random addresses associated with malicious software. The examiner has in no way addressed Applicants' previous arguments that Cheriton and Vaidya separately or together do not teach "processing said messages to form extrapolated connection sessions from said sampled data packets from which to determine a connection source that initiated the connection session." The examiner has not alleged that Jeffries provides such teachings, nor does the Jeffries patent application provide such teachings. Accordingly, even if the teachings of Jeffries could have been combined with the teachings of Cheriton and Vaidya as the examiner proposes, the claimed invention would not have resulted.

Jeffries does not teach processing messages from data sensors to form "extrapolated connection sessions from said sampled data packets from which to determine a connection source that initiated the connection session" as claimed. Instead of addressing this feature, the examiner alleges on page 3 of the Final Rejection that:

Jeffries teaches detecting the spread of malicious software (§21) by counting references of destination addresses (§§13-16) and comparing the count to a threshold (§16) to determine if a gateway may be receiving random addresses (which are associated with malicious software, §16). Therefore, it would have been obvious to one having ordinary skilled in the art at the time the invention was made to modify Cheriton to score each grouping of filter results based on

the quantity of attack destinations (destination addresses) and compare those values to a predetermined threshold.

Even if the above statements of the examiner were true, this is not what is claimed. As just noted, the claimed method and system form “extrapolated connection sessions” and detect a surveillance probe from scoring of the grouped connection sessions and comparing the scores to an empirically derived threshold. Simply counting unique destination addresses and taking hash values to determine if the destination addresses are non-random as suggested by Jeffries is not particularly relevant to what is claimed. As with Cheriton and Vaidya, the teachings of Jeffries are insufficient to properly distinguish scanning by malicious software or human activities. Applicant notes that there are many applications which, as a matter of course, make hundreds or thousands of connections in a short time for legitimate reasons. These include FTP data transfer connections due to mirroring, P2P applications, and browsers talking to portal web sites with many frames or images from many sites. The claimed method and system move beyond this by examining connections, which are very different from the flows of Cheriton or even bidirectional flows where the connection initiator is not known. Moreover, simply looking at packet types (such as RSTs, SYNs, and SYN-ACKs) to determine connection direction is error-prone either due to protocols that do not support such packet types (UDP) or due to sophisticated scanning techniques that use packet types in non-standard ways to attempt to evade security sensors. The claimed method and system avoid such issues by processing the messages from the data sensors to form extrapolated connection sessions from which good and bad connections may be identified based at least in part on the address of the connection source. The “bad” connections are then identified (*e.g.*, packets for which there are no replies or the replies are error messages) and counted to determine if thresholds suggestive of system surveillance preceding an attack are exceeded. Paragraphs [0057] and [0058] disclose exemplary techniques for detecting surveillance probes from the detected connections. Such teachings are not provided by Jeffries.

Accordingly, even if the teachings of Jeffries could have been combined with the teachings of Cheriton and Vaidya as the examiner alleges, neither the method of claim 3 nor the system of claim 22 would have been suggested to one skilled in the art. Withdrawal of the rejection of claims 3 and 28 is thus solicited.

Desai et al., Daizo, Trcka et al., and Ullmann et al. have been cited by the Examiner with respect to particular features of the dependent claims. Applicant submits that none of these patent documents teaches the afore-mentioned features of claims 3 and 22 that are not taught by Cheriton, Vaidya, or Jeffries. Accordingly, even if the teachings of one or more of these references would have been combined with the teachings of Cheriton, Vaidya, and Jeffries as the Examiner alleges, the features of claims 3 and 22 still would not have been suggested to one skilled in the art.

For at least the foregoing reasons, claim 3 and the claims dependent thereon (claims 5-12 and 19-21) are believed to be allowable over all of the cited references in any proposed combination. System claims 22-28 are believed to be allowable for the same reasons as claim 3. Allowance of claims 3, 5-12, and 19-28 is solicited.

Conclusion

Claims 3, 5-12, and 19-28 are believed to be novel and non-obvious over the cited references. Withdrawal of all rejections and issuance of a Notice of Allowability are solicited.

Date: September 17, 2009

/Michael P. Dunnam/
Michael P. Dunnam
Registration No. 32,611

Woodcock Washburn LLP
Cira Centre
2929 Arch Street, 12th Floor
Philadelphia, PA 19104-2891
Telephone: (215) 568-3100
Facsimile: (215) 568-3439